

# Well-Supported versus Approximate Nash Equilibria: Query Complexity of Large Games

Xi Chen\*

Columbia University

Yu Cheng†

University of Southern California

Bo Tang‡§

Oxford University

## Abstract

We study the randomized query complexity of approximate Nash equilibria (ANE) in large games. We prove that, for some constant  $\epsilon > 0$ , any randomized oracle algorithm that computes an  $\epsilon$ -ANE in a binary-action,  $n$ -player game must make  $2^{\Omega(n/\log n)}$  payoff queries. For the stronger solution concept of well-supported Nash equilibria (WSNE), Babichenko [Bab14] previously gave an exponential  $2^{\Omega(n)}$  lower bound for the randomized query complexity of  $\epsilon$ -WSNE, for some constant  $\epsilon > 0$ ; the same lower bound was shown to hold for  $\epsilon$ -ANE, but only when  $\epsilon = O(1/n)$ .

Our result answers an open problem posed by Hart and Nisan in [HN13] and by Babichenko in [Bab14], and is very close to the trivial upper bound of  $2^n$ . Our proof relies on a generic reduction from the problem of finding an  $\epsilon$ -WSNE to the problem of finding an  $\epsilon/(4\alpha)$ -ANE, in large games with  $\alpha$  actions, which might be of independent interest.

---

\*Email: [xichen@cs.columbia.edu](mailto:xichen@cs.columbia.edu). Supported in part by NSF grants CCF-1149257 and CCF-1423100.

†Email: [yu.cheng.1@usc.edu](mailto:yu.cheng.1@usc.edu). Supported in part by NSF grant CCF-1111270 and Shang-Hua Teng's Simons Investigator Award.

‡Email: [tangbonk1@gmail.com](mailto:tangbonk1@gmail.com). Supported in part by ERC grant 321171.

§This work was done in part while the authors were visiting the Simons Institute.

# 1 Introduction

The celebrated theorem of Nash [Nas50] states that every finite game has an equilibrium point. The solution concept of Nash equilibrium (NE) has been tremendously influential in economics and social sciences ever since (e.g. see [HR04]). The complexity and efficient approximation of NE have been studied intensively during the past decade, and much progress has been made (e.g., see [LMM03; AKV05; BVV05; KT07; TS07; DGP09; CDT09; EY10; Meh14; DP15; BPR15; CDO15; DKT15; Rub15; DKS15; Bar15]).

In this paper, we study the randomized query complexity of computing an  $\epsilon$ -approximate Nash equilibrium (ANE) in large games, for some *constant*  $\epsilon > 0$ . Given a game  $\mathcal{G}$  with  $n$  players and  $\alpha$  actions for each player, we index the players by the set  $[n] = \{1, \dots, n\}$  and index the actions by the set  $[\alpha] = \{1, \dots, \alpha\}$ . Recall that an  $\epsilon$ -ANE of  $\mathcal{G}$  is a mixed strategy profile

$$\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n), \quad \text{where } \mathbf{x}_i \in [0, 1]^\alpha \text{ sums to 1 for each } i \in [n],$$

in which each player  $i$  plays an  $\epsilon$ -best response  $\mathbf{x}_i$  to other players' strategies  $\mathbf{x}_{-i}$ <sup>1</sup> (see the precise definition in Section 2). Since the notion of ANE (as well as that of well-supported Nash equilibria to be discussed below) is additive, we always assume that the payoff functions of games considered throughout this paper take values between 0 and 1.

For the (payoff) query model, an oracle algorithm with unlimited computational power is given an approximation parameter  $\epsilon$ , the number of players  $n$  and the number of actions  $\alpha$  in an unknown game  $\mathcal{G}$ , and needs to find an  $\epsilon$ -ANE of  $\mathcal{G}$ . The algorithm has oracle access to the payoff functions of players in  $\mathcal{G}$ : For each round, the algorithm can adaptively query a pure strategy profile  $\mathbf{a} \in [\alpha]^n$ , and receives the payoff of every player with respect to  $\mathbf{a}$ . We are interested in the number of queries needed by any randomized oracle algorithm for this task. Note that a trivial upper bound is  $\alpha^n$ , by simply querying all the pure strategy profiles.

## Prior Results and Related Work

The query complexity of (approximate) Nash equilibria and related solution concepts has received considerable attention recently, e.g. see [FGGS13; HN13; FS14; GR14; Bab14; BB15; GT15]. Below we review results that are most relevant to our work.

The query complexity of (approximate) correlated equilibria<sup>2</sup> (CE) is well understood. For the (payoff) query model considered here, randomized algorithms exist (e.g., regret-minimizing algorithms [HMC00; Har05; BM07]) for finding an  $\epsilon$ -CE using  $\text{poly}(1/\epsilon, \alpha, n)$  many queries. It turns out that both randomization and approximation are necessary. Babichenko and Barman in [BB15] first showed that every deterministic algorithm that finds an exact CE requires exponentially many queries in  $n$ . Hart and Nisan [HN13] then showed that the same exponential lower bound holds for any deterministic algorithm for  $(1/2)$ -CE and any randomized algorithm for exact CE. Now for the stronger (expected payoff) query model, where the oracle returns the expected payoffs of any mixed strategy profile<sup>3</sup>, Papadimitriou and Roughgarden [PR08] and Jiang and Leyton-Brown [JL11] gave

<sup>1</sup>We follow the convention and write  $\mathbf{x}_{-i} := (\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_n)$ , strategies of players other than  $i$  in  $\mathbf{x}$ .

<sup>2</sup>An  $\epsilon$ -correlated equilibrium is a probability distribution over pure strategy profiles, i.e.  $[\alpha]^n$ , such that any player unilaterally deviating from strategies drawn from it can increase her expected payoff by no more than  $\epsilon$ .

<sup>3</sup>Such an oracle can be implemented in polynomial time for many classes of succinct games; see [PR08].

a deterministic algorithm that computes an exact CE in polynomial time using polynomially many queries (both in  $\alpha$  and  $n$ ).

Turning to the harder, but perhaps more interesting, problem of approximating Nash equilibria under the payoff query model, the deterministic lower bound of [HN13] for  $(1/2)$ -CE directly implies the same bound for  $(1/2)$ -ANE, since any  $\epsilon$ -ANE by definition is also an  $\epsilon$ -CE. For the randomized query complexity of NE, Babichenko [Bab14] showed that any randomized algorithm requires  $2^{\Omega(n)}$  queries to find an  $\epsilon$ -well-supported Nash equilibrium (WSNE), in a binary-action,  $n$ -player game (see [Theorem 1.4](#)). Recall that an  $\epsilon$ -WSNE of a game is a mixed strategy profile  $\mathbf{x}$  in which the probability of player  $i$  playing action  $j$  is positive only when action  $j$  is an  $\epsilon$ -best response with respect to  $\mathbf{x}_{-i}$  (see [Section 2](#) for the precise definition). By definition, an  $\epsilon$ -WSNE is also an  $\epsilon$ -ANE but the inverse is not true. Following a well-known connection between WSNE and ANE [DGP09] (and using random samples to approximate expected payoffs), Babichenko [Bab14] showed that the same  $2^{\Omega(n)}$  bound holds for the randomized query complexity of  $\epsilon$ -ANE, but only when  $\epsilon = O(1/n)$ . The randomized query complexity of  $\epsilon$ -ANE in large games, an arguably more natural relaxation of exact NE compared to WSNE, remains an open problem when  $\epsilon > 0$  is a constant.

## Our Results

For binary-action,  $n$ -player games, we show that  $2^{\Omega(n/\log n)}$  queries are required for any randomized algorithm to find an  $\epsilon$ -ANE, for some constant  $\epsilon > 0$ . To state the result, we use  $\text{QC}_p(\text{ANE}(n, \epsilon))$ , for some  $p > 0$ , to denote the smallest  $T$  such that there exists a randomized oracle algorithm that uses no more than  $T$  queries and outputs an  $\epsilon$ -ANE with probability at least  $p$ , given any unknown binary-action,  $n$ -player game. Our main result is the following lower bound on  $\text{QC}_p(\text{ANE}(n, \epsilon))$ :

**Theorem 1.1 (Main).** *There exist two constants  $\epsilon > 0$  and  $c > 0$  such that*

$$\text{QC}_p(\text{ANE}(n, \epsilon)) = 2^{\Omega(n/\log n)}, \quad \text{where } p = 2^{-cn/\log n}.$$

Our lower bound answers an open problem posed by Hart and Nisan [HN13] and by Babichenko [Bab14]. Our result shows that, in terms of their query complexities, finding an  $\epsilon$ -ANE is almost as hard as finding an  $\epsilon$ -WSNE in a large game, even for constant  $\epsilon > 0$ . It also implies the following corollary regarding the rate of convergence of  $k$ -queries dynamics (see [Bab14] for the definition).

**Corollary 1.2.** *There exist two constants  $\epsilon > 0$  and  $c > 0$  such that no  $k$ -queries dynamic can converge to an  $\epsilon$ -ANE in  $2^{\Omega(n/\log n)}/k$  steps with probability at least  $2^{-cn/\log n}$  in all binary-action and  $n$ -player games.*

In addition to the randomized query complexity, our proof of [Theorem 1.1](#) yields a polynomial-time reduction<sup>4</sup> from the problem of finding an  $\epsilon$ -WSNE to that of finding an  $(\epsilon' = \Omega(\epsilon))$ -ANE in a *succinct game* with a fixed number of actions. Following the definition from [PR08], we say that an  $\alpha$ -action succinct game is a pair  $(n, U)$ , where  $n$  is the number of players and  $U$  is a (multi-output) Boolean circuit that, given a pure strategy profile  $\mathbf{a} \in [\alpha]^n$  (encoded in binary), outputs the payoffs of all  $n$  players with respect to  $\mathbf{a}$  in the game. We show that

---

<sup>4</sup>Recall that a polynomial-time reduction from total search problem  $A$  to total search problem  $B$  is a pair  $(f, g)$  of polynomial-time computable functions such that: 1) for every input instance  $x$  of  $A$ ,  $f(x)$  is an input instance of  $B$ ; 2) for every solution  $y$  to  $f(x)$  in  $B$ ,  $g(y)$  is a solution to  $x$  in  $A$ .

**Theorem 1.3.** *Let  $\epsilon \geq 0$  and  $\alpha \in \mathbb{N}$  be two constants. Then the problem of finding an  $\epsilon$ -WSNE is polynomial-time reducible to that of finding an  $\epsilon/(4\alpha)$ -ANE, both in an  $\alpha$ -action succinct game.*

### Approximate vs Well-Supported Nash Equilibria

Let  $\text{QC}_p(\text{WSNE}(n, \epsilon))$ , for some  $p > 0$ , denote the smallest  $T$  such that there exists a randomized oracle algorithm that uses no more than  $T$  queries and outputs an  $\epsilon$ -WSNE with probability at least  $p$ , given any unknown binary-action,  $n$ -player game. Babichenko showed that

**Theorem 1.4** ([Bab14]). *There exist two constants  $\epsilon > 0$  and  $c > 0$  such that*

$$\text{QC}_p(\text{WSNE}(n, \epsilon)) = 2^{\Omega(n)}, \quad \text{where } p = 2^{-cn}.$$

Given the result of Babichenko as above, the same exponential lower bound for the randomized query complexity of  $\epsilon$ -ANE, for small enough constant  $\epsilon > 0$ , would follow immediately if

*Given oracle access to  $\mathcal{G}$  and any  $\epsilon'$ -ANE of  $\mathcal{G}$ , where  $\epsilon' = c(\alpha) \cdot \epsilon$  for some constant  $c > 0$  that only depends on  $\alpha$ , there is a query-efficient procedure that outputs an  $\epsilon$ -WSNE of  $\mathcal{G}$ .*

However, the best such procedure known is the following result from [DGP09]. The parameters are subsequently improved in [Bab14], where the number of queries needed is also analyzed:

*Given oracle access to  $\mathcal{G}$  and any  $\epsilon^2/(16n)$ -ANE of  $\mathcal{G}$ , there is a procedure that outputs an  $\epsilon$ -WSNE of  $\mathcal{G}$ , where  $n$  is the number of players, using  $\text{poly}(\alpha, n, 1/\epsilon)$  payoff queries.*

The procedure is very natural: For each player, reallocate probabilities on actions with a relatively low expected payoff to a best-response action. By [Theorem 1.4](#), such a procedure implies the same exponential lower bound for  $\epsilon$ -ANE [Bab14] but only when  $\epsilon$  is  $O(1/n)$ .

No better procedure is known. By definition, an ANE poses a slightly weaker condition on each player compared to that of a WSNE. More specifically, given mixed strategies of other players  $\mathbf{x}_{-i}$ , for an  $\epsilon$ -WSNE,  $\mathbf{x}_i$  must be supported on actions that are  $\epsilon$ -best responses to  $\mathbf{x}_{-i}$ , while in an  $\epsilon$ -ANE,  $\mathbf{x}_i$  can be any mixed strategy that yields an overall  $\epsilon$ -best response to  $\mathbf{x}_{-i}$ . For example,  $\mathbf{x}_i$  may put  $1 - \epsilon$  probability on best-response actions while putting  $\epsilon$  probability on any other actions. On the one hand, this makes WSNE much easier to analyze and control in hardness reductions, which is why it played a critical role in characterizing the complexity of Nash equilibria, starting with the work of [DGP09], later in [CDT09] and subsequent works. On the other hand, as the  $\epsilon$  being of interest in [DGP09; CDT09] is either exponentially or polynomially small, any hardness result for  $\epsilon$ -WSNE yields the same result for  $\epsilon$ -ANE (by combining the procedure of [DGP09] described above and a folklore padding argument).

### Our Approach

While we were not able to improve the procedure of [DGP09; Bab14], we prove [Theorem 1.1](#) via a *query-efficient* reduction from the problem of finding a WSNE to that of finding a ANE:

*Given any  $\alpha$ -action,  $n$ -player game  $\mathcal{G}$  and any parameter  $\epsilon > 0$ , one can define a new  $\alpha$ -action game  $\mathcal{G}'$  with a slightly larger set of  $O(\alpha^2 \log(n/\epsilon) \cdot n)$  players such that*

1. *To answer each payoff query on  $\mathcal{G}'$ , it suffices to make  $\alpha n$  payoff queries on  $\mathcal{G}$ ;*

2. There is a procedure that, given any  $\epsilon$ -ANE  $\mathbf{x}$  of  $\mathcal{G}'$ , outputs a  $(4\alpha\epsilon)$ -WSNE  $\mathbf{y}$  of  $\mathcal{G}$ , with no payoff oracle access to  $\mathcal{G}$  or  $\mathcal{G}'$ .

Our reduction is presented in [Section 3](#). [Theorem 1.1](#) then follows directly from the lower bound of Babichenko [[Bab14](#)] on the randomized query complexity of WSNE (in [Theorem 1.4](#)). [Theorem 1.3](#) follows from the fact that: 1) the payoff entries of  $\mathcal{G}'$  are easy to compute; and 2) the procedure to obtain  $\mathbf{y}$  from  $\mathbf{x}$  runs in time polynomial in the length of the binary representation of  $\mathbf{x}$ , when the number of actions  $\alpha$  is bounded.

Recall that in the procedure of [[DGP09](#); [Bab14](#)], an  $\epsilon$ -WSNE is obtained from an  $\epsilon'$ -ANE with  $\epsilon' = \epsilon^2/(16n)$  by reallocating probabilities on actions with relatively low expected payoff (formally, actions with payoff  $\Omega(\epsilon)$  lower than the best response) to best-response actions. From the definition of ANE, no player can have probability more than  $O(\epsilon'/\epsilon) = O(\epsilon/n)$  on actions with low payoff in an  $\epsilon'$ -ANE. Thus, the procedure changes the expected payoff of each player on each action by at most  $n \cdot O(\epsilon/n) = O(\epsilon)$  since it changes the mixed strategy of each player by  $O(\epsilon/n)$ . It follows that the new mixed strategy profile is an  $\epsilon$ -WSNE. The blow up of a factor of  $n$  from  $\epsilon'$  to  $\epsilon$  is precisely due to the cumulative impact on a player's expected payoff imposed by small changes to all other players' mixed strategies.

Our reduction from WSNE to ANE overcomes this obstacle by constructing a new and slightly larger game  $\mathcal{G}'$  with  $O(n \log n)$  players, where each player  $i$  in the original  $n$ -player game  $\mathcal{G}$  is now simulated by a group of  $O(\log n)$  players indexed by  $(i, j)$  in the new game  $\mathcal{G}'$ . The payoff function of player  $(i, j)$  in  $\mathcal{G}'$  is exactly the same as that of player  $i$  in  $\mathcal{G}$ , but is defined with respect to the *aggregate* action of each group of players in  $\mathcal{G}'$  by taking the *majority* among each group.

We then show that an  $\epsilon$ -WSNE of  $\mathcal{G}$  can be recovered from an  $\epsilon'$ -ANE of  $\mathcal{G}'$ , with  $\epsilon' = \Omega(\epsilon)$ , by 1) computing the distribution of the majority action of each group and 2) truncating the small entries in each distribution. Intuitively, by focusing on the aggregate behavior of each group of  $O(\log n)$  *independent* players in  $\mathcal{G}'$ , we make sure that the mixed strategies obtained from Step 1) are highly concentrated on actions with close-to-best expected payoffs, and actions with low payoffs can only appear as the majority action of a group with probability  $O(\epsilon/n)$ . Therefore in Step 2), we only need to truncate entries with probability  $O(\epsilon/n)$ , and the remaining positive entries would correspond to close-to-best actions. We can also control the effect of this truncation at the same time, because when the number of actions are bounded, the aggregate behavior of each group changes by at most  $O(\epsilon/n)$ . This allows us to show that the result is an  $\epsilon$ -WSNE of the original game  $\mathcal{G}$ .

## Organization

The rest of the paper is organized as follows. We first give formal definitions of ANE and WSNE in [Section 2](#). In [Section 3](#) we present the reduction from WSNE to ANE for large games, and then use it to prove [Theorem 1.1](#) and [Theorem 1.3](#) in [Section 4](#). We conclude and discuss open problems in [Section 5](#).

## 2 Preliminaries

A game  $\mathcal{G}$  is a triple  $(n, \alpha, \mathbf{u})$ , where  $n$  is the number of players,  $\alpha$  is the number of actions for each player, and  $\mathbf{u} = (u_1, \dots, u_n)$  are the payoff functions, one for each player. We always use  $[n] = \{1, \dots, n\}$  to denote the set of players and  $[\alpha] = \{1, \dots, \alpha\}$  to denote the set of actions for each player. Since we are interested in additive approximations, each  $u_i$  maps  $[\alpha]^n$  to  $[0, 1]$ .

Let  $\Delta_\alpha$  denote the set of probability distributions over  $[\alpha]$ . A mixed strategy profile of  $\mathcal{G}$  is then a tuple  $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  of mixed strategies, where  $\mathbf{x}_i \in \Delta_\alpha$  denotes the mixed strategy of player  $i$ . Given  $\mathbf{x}$ , we use  $\mathbf{x}_{-i}$  to denote the tuple of mixed strategies of all players other than  $i$ . As a shorthand, we write  $u_i(\mathbf{x})$  to denote the expected payoff of player  $i$  with respect to  $\mathbf{x}$ , and write  $u_i(a, \mathbf{x}_{-i})$  to denote the expected payoff of player  $i$  playing action  $a \in [\alpha]$  with respect to  $\mathbf{x}_{-i}$ :

$$u_i(\mathbf{x}) = \mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] \quad \text{and} \quad u_i(a, \mathbf{x}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{x}_{-i}} [u_i(a, \mathbf{b})].$$

Next we define approximate and well-supported Nash equilibria.

**Definition 1.** Given  $\epsilon > 0$ , an  $\epsilon$ -approximate Nash equilibrium of an  $\alpha$ -action and  $n$ -player game  $\mathcal{G}(n, \alpha, \mathbf{u})$  is a mixed strategy profile  $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  such that for every player  $i \in [n]$ :

$$u_i(\mathbf{x}) \geq u_i(a', \mathbf{x}_{-i}) - \epsilon, \quad \text{for all } a' \in [\alpha].$$

**Definition 2.** Given  $\epsilon > 0$ , an  $\epsilon$ -well-supported Nash equilibrium of  $\mathcal{G}(n, \alpha, \mathbf{u})$  is a mixed strategy profile  $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  such that for every player  $i \in [n]$  and every action  $a$  in the support of  $\mathbf{x}_i$ :

$$u_i(a, \mathbf{x}_{-i}) \geq u_i(a', \mathbf{x}_{-i}) - \epsilon, \quad \text{for all } a' \in [\alpha].$$

Finally, we give a formal definition of succinct games [PR08].

**Definition 3.** An  $\alpha$ -action succinct game is a pair  $(n, U)$ , where  $n$  is the number of players and  $U$  is a (multi-output) Boolean circuit that, given any pure strategy profile  $\mathbf{a} \in [\alpha]^n$  (encoded in binary), outputs the payoffs of all  $n$  players with respect to  $\mathbf{a}$  in the game. The input size of  $(n, U)$  is the size of the circuit  $U$ .

### 3 A Reduction from Well-Supported to Approximate Nash Equilibria

Given an  $\alpha$ -action,  $n$ -player game  $\mathcal{G}(n, \alpha, \mathbf{u})$  and  $\epsilon \in (0, 1)$ , we define a new game  $\mathcal{G}'(sn, \alpha, \mathbf{u}')$  with  $sn$  players, where

$$s = 2\alpha^2 \cdot \lceil \ln(n/\epsilon) \rceil.$$

We then show that, given an  $\epsilon$ -ANE  $\mathbf{x}$  of the new game  $\mathcal{G}'$ , one can compute a  $(4\alpha\epsilon)$ -WSNE  $\mathbf{y}$  of  $\mathcal{G}$  without making any payoff queries to  $\mathcal{G}$  or  $\mathcal{G}'$ .

For each player  $i \in [n]$  in  $\mathcal{G}$ , we introduce a group of  $s$  players in  $\mathcal{G}'$ , indexed by  $(i, j)$  with  $j \in [s]$ , and use  $u'_{i,j}$  to denote the payoff function of player  $(i, j)$ . Given any pure strategy profile  $\mathbf{a} = (a_{i,j} : i \in [n], j \in [s])$ , we define the payoff  $u'_{i,j}(\mathbf{a})$  of player  $(i, j)$  as follows. First, for each  $i \in [n]$ , let  $\bar{a}_i \in [\alpha]$  denote the *majority* action played by the  $i$ -th group (players  $(i, j)$ ,  $j \in [s]$ ) in the pure strategy profile  $\mathbf{a}$  (break ties by choosing the action with the smallest index). Write  $\bar{\mathbf{a}} = (\bar{a}_1, \dots, \bar{a}_n)$ . Next, the payoff of player  $(i, j)$  under  $\mathbf{a}$  is defined as

$$u'_{i,j}(\mathbf{a}) = u_i(a_{i,j}, \bar{\mathbf{a}}_{-i}). \tag{1}$$

This completes the definition of  $\mathcal{G}'$ . The lemma below follows directly from the definition.

**Lemma 3.1.** *To answer a payoff query on  $\mathcal{G}'$ , it suffices to make  $\alpha n$  payoff queries on  $\mathcal{G}$ .*

*Proof.* By the definition of  $\mathcal{G}'$ ,  $u'_{i,j}(\mathbf{a})$ 's for all  $(i, j)$ ,  $i \in [n]$  and  $j \in [s]$ , are determined by

$$(u_i(a', \bar{\mathbf{a}}_{-i}) : i \in [n], a' \in [\alpha]),$$

for which  $\alpha n$  payoff queries on  $\mathcal{G}$  suffice.  $\square$

We conclude our reduction by proving the following lemma:

**Lemma 3.2.** *Given any  $\epsilon$ -ANE  $\mathbf{x}$  of  $\mathcal{G}'$ , one can compute a  $(4\alpha\epsilon)$ -WSNE  $\mathbf{y}$  of  $G$  without making any payoff queries on  $\mathcal{G}$  or  $\mathcal{G}'$ . Moreover, when  $\alpha$  is a constant, the computation of  $\mathbf{y}$  from  $\mathbf{x}$  can be done in time polynomial in the number of bits needed in the binary representation of  $\mathbf{x}$  and  $1/\epsilon$ .*

*Proof.* Let  $\mathbf{x} = (\mathbf{x}_{i,j} : i \in [n], j \in [\alpha])$  be an  $\epsilon$ -ANE of  $\mathcal{G}'$ . For each group  $i$  and action  $k \in [\alpha]$ , let

$$\bar{x}_{i,k} = \Pr_{\mathbf{a} \sim \mathbf{x}} [\bar{a}_i = k]. \quad (2)$$

Recall that  $\bar{a}_i$  is the majority action played by players  $(i, j)$ ,  $j \in [s]$ , in the pure strategy profile  $\mathbf{a}$ . Then by definition, each  $\bar{\mathbf{x}}_i = (\bar{x}_{i,1}, \dots, \bar{x}_{i,\alpha})$  is a probability distribution over  $[\alpha]$ .

Next we define a mixed strategy  $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$  of  $\mathcal{G}$ , and show that  $\mathbf{y}$  is a  $(4\alpha\epsilon)$ -WSNE. Let

$$c_{i,k} = \begin{cases} \bar{x}_{i,k} & \text{if } \bar{x}_{i,k} \leq \epsilon/n \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad y_{i,k} = \frac{\bar{x}_{i,k} - c_{i,k}}{1 - \sum_{j \in [\alpha]} c_{i,j}}. \quad (3)$$

It is clear that  $\mathbf{y}_i = (y_{i,1}, \dots, y_{i,\alpha})$  is a probability distribution over  $[\alpha]$ .

Now assume for contradiction that  $\mathbf{y}$  is not a  $(4\alpha\epsilon)$ -WSNE. Then for some player  $i \in [n]$ , there exists an action  $\ell \in [\alpha]$  such that  $y_{i,\ell} > 0$  but

$$\max_{k \in [\alpha]} u_i(k, \mathbf{y}_{-i}) > u_i(\ell, \mathbf{y}_{-i}) + 4\alpha\epsilon. \quad (4)$$

But note that, the total variation distance between  $\bar{\mathbf{x}}_j$  and  $\mathbf{y}_j$  for each  $j \in [n]$  is at most  $\alpha\epsilon/n$ . So by coupling and applying union bound, we have that

$$|u_i(k, \bar{\mathbf{x}}_{-i}) - u_i(k, \mathbf{y}_{-i})| \leq (n-1) \cdot (\alpha\epsilon/n) < \alpha\epsilon, \quad \text{for all } k \in [\alpha]. \quad (5)$$

It then follows from (4) and (5) that

$$\max_{k \in [\alpha]} u_i(k, \bar{\mathbf{x}}_{-i}) > u_i(\ell, \bar{\mathbf{x}}_{-i}) + 2\alpha\epsilon. \quad (6)$$

By the definition (1) of the payoff function  $u'_{i,j}$ , we have

$$u'_{i,j}(k, \mathbf{x}_{-i}) = u_i(k, \bar{\mathbf{x}}_{-i}), \quad \text{for all } j \in [s] \text{ and } k \in [\alpha]. \quad (7)$$

Combining (6) and (7), we have that for every player  $(i, j)$ ,  $j \in [s]$ :

$$\max_{k \in [\alpha]} u'_{i,j}(k, \mathbf{x}_{-i}) - u'_{i,j}(\ell, \mathbf{x}_{-i}) \geq 2\alpha\epsilon.$$



Since  $\mathbf{x}$  is an  $\epsilon$ -ANE of  $\mathcal{G}'$ ,  $x_{i,j,\ell} \leq 1/(2\alpha)$ . By Hoeffding bound and plugging in  $s = 2\alpha^2 \cdot \lceil \ln(n/\epsilon) \rceil$ ,

$$\begin{aligned}\bar{x}_{i,\ell} &= \Pr[\ell \text{ is the majority action among players } (i, j), j \in [s]] \\ &\leq \Pr[\text{the number of players } (i, j) \text{ playing } \ell \text{ is at least } s/\alpha] \leq e^{-s/(2\alpha^2)} \leq \epsilon/n.\end{aligned}$$

By (3), this implies that  $y_{i,\ell} = 0$ , which contradicts our assumption and proves that  $\mathbf{y}$  is indeed a  $(4\alpha\epsilon)$ -WSNE of  $\mathcal{G}$ . From the definition of  $\mathbf{y}$ , it is clear that the computation of  $\mathbf{y}$  from  $\mathbf{x}$  does not require any payoff queries.

For the running time, when  $\alpha$  is a constant, to compute  $\bar{x}_{i,k}$  in (2) one needs to go through

$$\alpha^s = \alpha^{2\alpha^2 \cdot \lceil \ln(n/\epsilon) \rceil} = (n/\epsilon)^{O(1)}$$

many pure strategy profiles of players  $(i, j)$ ,  $j \in [s]$ . Thus  $\mathbf{y}$  can be computed in time polynomial in the number of bits needed in the binary representation of  $\mathbf{x}$  and  $1/\epsilon$ .  $\square$

## 4 Proofs of Theorems 1.1 and 1.3

We use the query-efficient reduction given above to prove Theorem 1.1 and Theorem 1.3.

*Proof of Theorem 1.1.* By Theorem 1.4, there exist two constants  $\epsilon' > 0$  and  $c' > 0$  such that

$$\text{QC}_{p'}(\mathbf{WSNE}(n', \epsilon)) = 2^{\Omega(n')}, \quad \text{where } p' = 2^{-c'n'}.$$

Let  $n = 8n' \cdot \lceil \ln(n'/\epsilon') \rceil$  and  $\epsilon = 8\epsilon'$ . It follows from Lemma 3.1 and Lemma 3.2 that

$$\text{QC}_{p'}(\mathbf{ANE}(n, \epsilon)) \geq \text{QC}_{p'}(\mathbf{WSNE}(n', \epsilon)) = 2^{\Omega(n')}.$$

The theorem then follows from  $n' = \Omega(n/\log n)$ .  $\square$

*Proof of Theorem 1.3.* From Lemma 3.2, it suffices to show that, given any  $\alpha$ -action succinct game  $\mathcal{G} = (n, U)$ , one can construct, in polynomial time, a Boolean circuit  $U'$  that implements the payoff functions of players in  $\mathcal{G}'$ . This can be done by following the definition of  $\mathcal{G}'$  in the previous section, as the payoffs of a pure strategy profile  $\mathbf{a}$  in  $\mathcal{G}'$  only depends (in a straight-forward fashion) on the payoffs of  $\alpha n = O(n)$  easy-to-compute profiles of  $\mathcal{G}$ .  $\square$

## 5 Conclusion

In this paper, we present a simple and efficient reduction from the problem of finding a WSNE to that of finding an ANE in large games with a bounded number of actions. Our results complement the existing study on relations between WSNE and ANE. As an application, we obtain a lower bound on the randomized query complexity of  $\epsilon$ -ANE for some constant  $\epsilon > 0$ . It would be interesting to see other applications of our reduction in understanding the complexity of Nash equilibria. It also remains an open problem to remove the  $\log n$  factor in the exponent of our lower bound, i.e. to show that the number of queries needed to reach an  $\epsilon$ -ANE is indeed  $2^{\Omega(n)}$ . This  $\log n$  factor shows up because we simulate each player in the original game with  $O(\log n)$  players in the new game. Is there a more efficient simulation that uses fewer players?



## References

- [AKV05] T. Abbott, D. Kane, and P. Valiant. “On the Complexity of Two-Player Win-Lose Games”. In: *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. 2005, pp. 113–122 (cit. on p. 2).
- [Bab14] Y. Babichenko. “Query Complexity of Approximate Nash Equilibria”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. 2014, pp. 535–544 (cit. on pp. 1–5).
- [Bar15] S. Barman. “Approximating Nash Equilibria and Dense Bipartite Subgraphs via an Approximate Version of Caratheodory’s Theorem”. In: *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*. 2015, pp. 361–369 (cit. on p. 2).
- [BB15] Y. Babichenko and S. Barman. “Query Complexity of Correlated Equilibrium”. In: *ACM Transactions on Economics and Computation* 3.4 (2015), 22:1–22:9 (cit. on p. 2).
- [BM07] A. Blum and Y. Monsoir. *From External to Internal Regret*. University of Chicago Press, 2007 (cit. on p. 2).
- [BPR15] Y. Babichenko, C.H. Papadimitriou, and A. Rubinstein. “Can Almost Everybody be Almost Happy? PCP for PPAD and the Inapproximability of Nash”. In: *CoRR* abs/1504.02411 (2015) (cit. on p. 2).
- [BVV05] I. Bárány, S. Vempala, and A. Vetta. “Nash Equilibria in Random Games”. In: *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. 2005, pp. 123–131 (cit. on p. 2).
- [CDO15] X. Chen, D. Durfee, and A. Orfanou. “On the Complexity of Nash Equilibria in Anonymous Games”. In: *Proceedings of the 46th ACM Symposium on Theory of Computing*. 2015, pp. 381–390 (cit. on p. 2).
- [CDT09] X. Chen, X. Deng, and S.-H. Teng. “Settling the Complexity of Computing Two-player Nash Equilibria”. In: *Journal of the ACM* 56.3 (2009), pp. 1–57 (cit. on pp. 2, 4).
- [DGP09] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. “The Complexity of Computing a Nash Equilibrium”. In: *SIAM Journal on Computing* 39.1 (2009) (cit. on pp. 2–5).
- [DKS15] I. Diakonikolas, D.M. Kane, and A. Stewart. “Nearly Optimal Learning and Sparse Covers for Sums of Independent Integer Random Variables”. In: *CoRR* abs/1505.00662 (2015) (cit. on p. 2).
- [DKT15] C. Daskalakis, G. Kamath, and C. Tzamos. “On the Structure, Covering, and Learning of Poisson Multinomial Distributions”. In: *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science*. 2015 (cit. on p. 2).
- [DP15] C. Daskalakis and C.H. Papadimitriou. “Approximate Nash Equilibria in Anonymous Games”. In: *Journal of Economic Theory* 156 (2015), pp. 207–245 (cit. on p. 2).
- [EY10] K. Etessami and M. Yannakakis. “On the Complexity of Nash Equilibria and Other Fixed Points”. In: *SIAM Journal on Computing* 39.6 (2010), pp. 2531–2597 (cit. on p. 2).

- [FGGS13] J. Fearnley, M. Gairing, P. Goldberg, and R. Savani. “Learning Equilibria of Games via Payoff Queries”. In: *Proceedings of the 14th ACM Conference on Electronic Commerce*. 2013, pp. 397–414 (cit. on p. 2).
- [FS14] J. Fearnley and R. Savani. “Finding Approximate Nash Equilibria of Bimatrix Games via Payoff Queries”. In: *Proceedings of the 15th ACM Conference on Economics and Computation*. 2014, pp. 657–674 (cit. on p. 2).
- [GR14] P.W. Goldberg and A. Roth. “Bounds for the Query Complexity of Approximate Equilibria”. In: *Proceedings of the 15th ACM Conference on Economics and Computation*. 2014, pp. 639–656 (cit. on p. 2).
- [GT15] P.W. Goldberg and S. Turchetta. “Query Complexity of Approximate Equilibria in Anonymous Games”. In: *Proceedings of the 11th Conference on Web and Internet Economics*. 2015 (cit. on p. 2).
- [Har05] S. Hart. “Adaptive Heuristics”. In: *Econometrica* 73.5 (2005), pp. 1401–1430 (cit. on p. 2).
- [HMC00] S. Hart and A. Mas-Colell. “A Simple Adaptive Procedure Leading to Correlated Equilibrium”. In: *Econometrica* 68.5 (2000), pp. 1127–1150 (cit. on p. 2).
- [HN13] S. Hart and N. Nisan. “The Query Complexity of Correlated Equilibria”. In: *CoRR* abs/1305.4874 (2013) (cit. on pp. 1–3).
- [HR04] C.A. Holt and A.E. Roth. “The Nash Equilibrium: A Perspective”. In: *Proceedings of the National Academy of Sciences* 101.12 (2004), pp. 3999–4002 (cit. on p. 2).
- [JL11] A.X. Jiang and K. Leyton-Brown. “Polynomial-time Computation of Exact Correlated Equilibrium in Compact Games”. In: *Proceedings of the 12th ACM Conference on Electronic Commerce*. 2011 (cit. on p. 2).
- [KT07] R. Kannan and T. Theobald. “Games of Fixed Rank: A Hierarchy of Bimatrix Games”. In: *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*. 2007, pp. 1124–1132 (cit. on p. 2).
- [LMM03] R.J. Lipton, E. Markakis, and A. Mehta. “Playing Large Games Using Simple Strategies”. In: *Proceedings of the 4th ACM Conference on Electronic Commerce*. 2003, pp. 36–41 (cit. on p. 2).
- [Meh14] R. Mehta. “Constant Rank Bimatrix Games Are PPAD-hard”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. 2014, pp. 545–554 (cit. on p. 2).
- [Nas50] J.F. Nash. “Equilibrium points in N-Person Games”. In: *Proceedings of the National Academy of Sciences* 36.1 (1950), pp. 48–49 (cit. on p. 2).
- [PR08] C.H. Papadimitriou and T. Roughgarden. “Computing Correlated Equilibria in Multi-Player Games”. In: *Journal of the ACM* 55.3 (2008) (cit. on pp. 2, 3, 6).
- [Rub15] A. Rubinstein. “Inapproximability of Nash Equilibrium”. In: *Proceedings of the 46th ACM Symposium on Theory of Computing*. 2015, pp. 409–418 (cit. on p. 2).
- [TS07] H. Tsaknakis and P.G. Spirakis. “An Optimization Approach for Approximate Nash Equilibria”. In: *Proceedings of the 3rd International Workshop on Internet and Network Economics*. 2007, pp. 42–56 (cit. on p. 2).